



Intel[®] Pentium[®] M Processor

Specification Update

November 2005

Notice: The Intel[®] Pentium[®] M processor may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are documented in this Specification Update.

Document Number: **252665-017**



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Intel® Pentium® M processor may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

[†]Hyper-Threading Technology requires a computer system with a Mobile Intel Pentium 4 Processor, a chipset and BIOS that utilize this technology, and an operating system that includes optimizations for this technology. Performance will vary depending on the specific hardware and software you use. See <http://www.intel.com/info/hyperthreading> for information.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Intel, Pentium, Celeron, Intel Xeon, Intel SpeedStep, MMX and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2003 – 2005, Intel Corporation. All rights reserved.



Contents

Preface	5
Summary Tables of Changes	7
Identification Information	11
Errata	13
Specification Changes	26
Specification Clarifications	27
Documentation Changes	30

Revision History

Revision Number	Description	Date
001	Initial Release	March 2003
002	<ul style="list-style-type: none"> Added Erratum Y15 Clarified Status description for all Erratum 	April 2003
003	<ul style="list-style-type: none"> Updated Pentium M processor Identification Table 	June 2003
004	<ul style="list-style-type: none"> Added Erratum Y16, Y17 	July 2003
005	<ul style="list-style-type: none"> Added Erratum Y18 	September 2003
006	<ul style="list-style-type: none"> Added Erratum Y19 - Y21 	November 2003
007	<ul style="list-style-type: none"> Updated Erratum Y20 Added Erratum Y22 	December 2003
008	<ul style="list-style-type: none"> Added Erratum Y23 and Y24 Added Specification Clarification Y1 Updated Processor Identification table 	April 2004
009	<ul style="list-style-type: none"> Added Erratum Y25 and Y26 	May 2004
010	<ul style="list-style-type: none"> Added Erratum Y27, Y28, Y29 and Y30 	October 2004
011	<ul style="list-style-type: none"> Added Erratum Y31, Y32 and Y33 	November 2004
012	<ul style="list-style-type: none"> Added Erratum Y34 and Y35 	December 2004
013	<ul style="list-style-type: none"> Updated Summary Tables of Changes Added Erratum Y36 	February 2005
014	<ul style="list-style-type: none"> Added Erratum Y37 Added Specification Clarification Y1 	May 2005
015	<ul style="list-style-type: none"> Updated Summary Tables of Changes Removed Erratum Y28 – Y30 (which were duplicates of Y4 – Y6). Added Erratum Y38 – Y40 	June 2005
016	<ul style="list-style-type: none"> Updated Affected and Related Documents Tables 	September 2005
017	<ul style="list-style-type: none"> Added Erratum Y41 	November 2005

§



Preface

This document is an update to the specifications contained in the documents listed in the following Affected Documents/Related Documents table. It is a compilation of device and document errata and specification clarifications and changes, and is intended for hardware system manufacturers and for software developers of applications, operating system, and tools.

Information types defined in the Nomenclature section of this document are consolidated into this update document and are no longer published in other documents. This document may also contain information that has not been previously published.

Affected Documents

Document Title	Document Number
<i>Intel® Pentium® M Processor Datasheet</i>	252612-003

Related Documents

Document Title	Document Number
<i>IA-32 Intel® Architecture Software Developer's Manual, Volume 1: Basic Architecture</i>	253665
<i>IA-32 Intel® Architecture Software Developer's Manual, Volume 2A: Instruction Set Reference, A-M</i>	253666
<i>IA-32 Intel® Architecture Software Developer's Manual, Volume 2B: Instruction Set Reference, N-Z</i>	253667
<i>IA-32 Intel® Architecture Software Developer's Manual, Volume 3: System Programming Guide</i>	253668
<i>IA-32 Intel® Architecture Optimization Reference Manual</i>	248966

Nomenclature

S-Spec Number is a five-digit code used to identify products. Products are differentiated by their unique characteristics, e.g., core speed, L2 cache size, package type, etc. as described in the processor identification information table. Care should be taken to read all notes associated with each S-Spec number.

Errata are design defects or errors. Errata may cause the Intel® Pentium® M processor's behavior to deviate from published specifications. Hardware and software, designed to be used with any given processor, must assume that all errata documented for that processor are present on all devices unless otherwise noted.

Documentation Changes include typos, errors, or omissions from the current published specifications. These changes will be incorporated in the next release of the specifications.

Specification Clarifications describe a specification in greater detail or further highlight a specification's impact to a complex design situation. These clarifications will be incorporated in the next release of the specifications.

Specification Changes are modifications to the current published specifications for the Pentium M processor. These changes will be incorporated in the next release of the specifications.



Summary Tables of Changes

The following table indicates the Errata, Documentation Changes, Specification Clarifications, or Specification Changes that apply to Pentium M processors. Intel intends to fix some of the errata in a future stepping of the component and to account for the other outstanding issues through documentation or specification changes as noted. This table uses the following notations:

Codes Used in Summary Table

Stepping

X: Erratum, Specification Change or Clarification that applies to this stepping.

(No mark) or (Blank Box): This erratum is fixed in listed stepping or specification change does not apply to listed stepping.

Status

Doc: Document change or update that will be implemented.

PlanFix: This erratum may be fixed in a future of the product.

Fixed: This erratum has been previously fixed.

NoFix: There are no plans to fix this erratum.

Shaded: This item is either new or modified from the previous version of the document.

Each Specification Update item is prefixed with a capital letter to distinguish the product. The key below details the letters that are used in Intel's microprocessor Specification Updates:

A = Intel® Pentium® II processor
B = Mobile Intel® Pentium® II processor
C = Intel® Celeron® processor
D = Intel® Pentium® II Xeon™ processor
E = Intel® Pentium® III processor
F = Intel® Pentium® processor Extreme Edition and Intel® Pentium® D processor
G = Intel® Pentium® III Xeon™ processor
H = Mobile Intel® Celeron® processor at 466 MHz, 433 MHz, 400 MHz, 366 MHz, 333 MHz, 300 MHz, and 266 MHz
J = 64-bit Intel® Xeon™ processor MP with 1 MB L2 Cache

K = Mobile Intel® Pentium® III Processor – M
 L = Intel® Celeron® D processor
 M = Mobile Intel® Celeron® processor
 N = Intel® Pentium® 4 processor
 O = Intel® Xeon™ processor MP
 P = Intel® Xeon™ processor
 Q = Mobile Intel® Pentium® 4 processor supporting Hyper-Threading Technology[†] on 90-nm technology process
 R = Intel® Pentium® 4 processor on 90 nm process
 S = 64-bit Intel® Xeon™ processor with 800 MHz system bus (1 MB and 2 MB L2 cache versions)
 T = Mobile Intel® Pentium® 4 processor – M
 U = 64-bit Intel® Xeon™ processor MP with up to 8 MB L3 Cache
 V = Mobile Intel® Celeron® processor on .13 Micron process in Micro-FCPGA Package
 W = Intel® Celeron® M processor
 X = Intel® Pentium® M processor on 90 nm process with 2-MB L2 Cache
 Y = Intel® Pentium® M processor
 Z = Mobile Intel® Pentium® 4 Processor with 533 MHz System Bus

Note: The Specification Updates for the Pentium® processor, Pentium® Pro processor, and other Intel products do not use this convention.

NO.	Stepping	Plans	ERRATA
	B1		
Y1	X	NoFix	Performance Monitoring Event That Counts Intel® Thermal Monitor 2 Transitions (59h) Is Not Accurate
Y2	X	NoFix	Performance Monitoring Event That Counts the Number of Instructions Decoded (D0h) Is Not Accurate
Y3	X	NoFix	RDTSC Instruction May Report the Wrong Time-Stamp Counter Value
Y4	X	NoFix	Code Segment Limit Violation May Occur on 4-Byte Limit Check
Y5	X	NoFix	FST Instruction with Numeric and Null Segment Exceptions May Cause General Protection Faults to Be Missed and FP Linear Address (FLA) Mismatch
Y6	X	NoFix	Code Segment (CS) Is Wrong on SMM Handler When SMBASE Is Not Aligned
Y7	X	NoFix	IFU/BSU Deadlock May Cause System Hang
Y8	X	NoFix	Processor Can Enter a Livelock Condition under Certain Conditions When FP Exception Is Pending
Y9	X	NoFix	Write Cycle of Write Combining Memory Type Does Not Self Snoop
Y10	X	NoFix	Performance Monitoring Event That Counts Floating Point Computational Exceptions (11h) Is Not Accurate.
Y11	X	NoFix	Inconsistent Reporting of Data Breakpoints on FP (Intel® MMX technology) Loads
Y12	X	NoFix	Code Breakpoint May Be Taken after POP SS instruction If It Is followed by an Instruction That Faults
Y13	X	NoFix	SysEnter and SysExit Instructions May Write Incorrect Requestor Privilege Level (RPL) in the FP Code Segment Selector (FCS)

NO.	Stepping	Plans	ERRATA
	B1		
Y14	X	NoFix	Memory Aliasing with Inconsistent A and D Bits May Cause Processor Deadlock
Y15	X	NoFix	RDMSR or WRMSR to Invalid MSR Address May Not Cause GP Fault
Y16	X	NoFix	FP Tag Word Corruption
Y17	X	NoFix	Unable to Disable Reads/Writes to Performance Monitoring Related MSRs
Y18	X	NoFix	Move to Control Register Instruction May Generate a Breakpoint Report
Y19	X	NoFix	REP MOVSB Operation in Fast String Mode Continues in That Mode When Crossing Into a Page with a Different Memory Type
Y20	X	NoFix	The FXSAVE, STOS, or MOVSB Instruction May Cause a Store Ordering Violation When Data Crosses a Page with a UC Memory Type
Y21	X	NoFix	Machine Check Exception May Occur Due to Improper Line Eviction in the IFU
Y22	X	NoFix	POPF and POPFD Instructions That Set the Trap Flag Bit May Cause Unpredictable Processor Behavior
Y23	X	NoFix	Performance Event Counter Returns Incorrect Value on L2_LINES_IN Event
Y24	X	NoFix	VM Bit Will Be Cleared on a Double Fault Handler
Y25	X	NoFix	Code Fetch Matching Disabled Debug Register May Cause Debug Exception
Y26	X	NoFix	Upper Four PAT Entries Not Usable with Mode B or Mode C Paging
Y27	X	NoFix	SSE/SSE2 Streaming Store Resulting in a Self-Modifying Code (SMC) Event May Cause Unexpected Behavior
Y28			Removed, see Erratum Y4
Y29			Removed, see Erratum Y5
Y30			Removed, see Erratum Y6
Y31	X	NoFix	Page with PAT (Page Attribute Table) Set to USWC (Uncacheable Speculative Write Combine) While Associated MTRR (Memory Type Range Register) Is UC (Uncacheable) May Consolidate to UC
Y32	X	NoFix	Under Certain Conditions LTR (Load Task Register) Instruction May Result in System Hang
Y33	X	NoFix	Loading from Memory Type USWC (Uncacheable Speculative Write Combine) May Get Its Data Internally Forwarded from a Previous Pending Store
Y34	X	NoFix	FXSAVE after FNINIT without an Intervening FP (Floating Point) Instruction May Save Uninitialized Values for FDP (x87 FPU Instruction Operand (Data) Pointer Offset) and FDS (x87 FPU Instruction Operand (Data) Pointer Selector)
Y35	X	NoFix	FSTP (Floating Point Store) Instruction Under Certain Conditions May Result In Erroneously Setting a Valid Bit on an FP (Floating Point) Stack Register
Y36	X	NoFix	Snoops during the Execution of a HLT (Halt) Instruction May Lead to Unpredictable System Behavior
Y37	X	NoFix	Invalid Entries in Page-Directory-Pointer-Table Register (PDPTR) May Cause General Protection (#GP) Exception if the Reserved Bits are Set to One
Y38	X	NoFix	INIT Does Not Clear Global Entries in the TLB
Y39	X	NoFix	Use of Memory Aliasing with Inconsistent Memory Type May Cause System Hang

NO.	Stepping	Plans	ERRATA
	B1		
Y40	X	NoFix	Machine Check Exception May Occur When Interleaving Code between Different Memory Types
Y41	X	NoFix	Split I/O Writes Adjacent to Retry of APIC End of Interrupt (EOI) Request May Cause Livelock Condition

Number	SPECIFICATION CHANGE
	There are no Specification Changes in this Specification Update revision.

Number	SPECIFICATION CLARIFICATIONS
Y1	Specification Clarification with Respect to Time-stamp Counter

Number	DOCUMENTATION CHANGES
	There are no Documentation Changes in this Specification Update revision.

Identification Information

The Pentium M processor can be identified by the following values:

Family ¹	Model ²	Brand ID ³
0110	1001	00010110

NOTES:

1. The Family corresponds to bits [11:8] of the EDX register after Reset, bits [11:8] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register.
2. The Model corresponds to bits [7:4] of the EDX register after Reset, bits [7:4] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register.
3. The Brand ID corresponds to bits [7:0] of the EBX register after the CPUID instruction is executed with a1 in the EAX register.

Table 1. Intel® Pentium® M Processor Identification

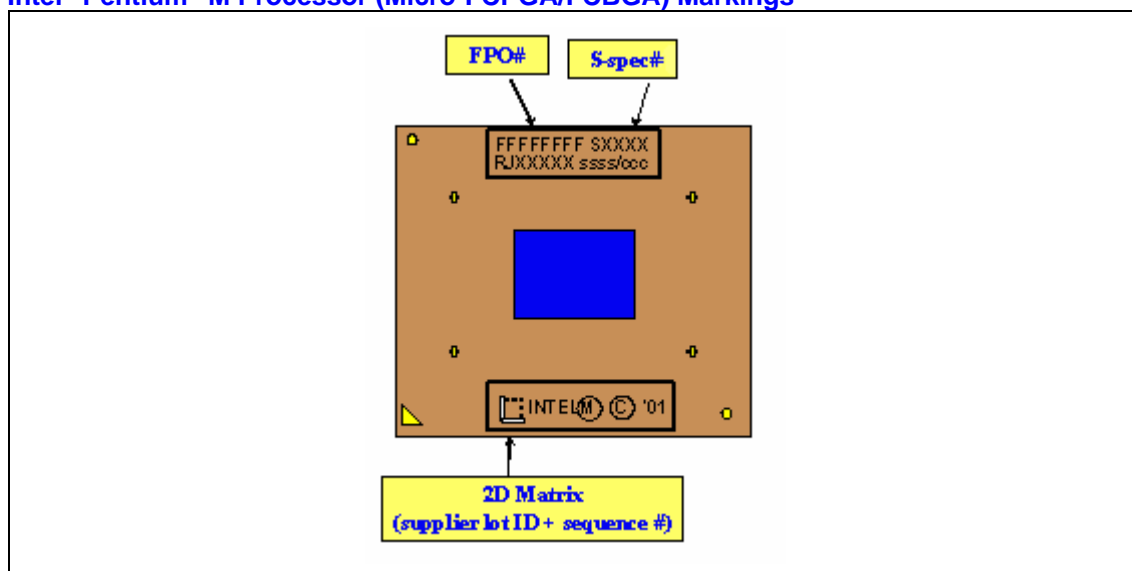
S-Spec/ QDF	Product Stepping	CPUID	Core Speed		Bus Frequency	Package	Notes
			Highest Frequency Mode (HFM)	Lowest Frequency Mode (LFM)			
SL6N4	B-1	0695h	1.30 GHz	600 MHz	400 MHz	Micro-FCPGA	2,6
SL6F8	B-1	0695h	1.40 GHz	600 MHz	400 MHz	Micro-FCPGA	2,5
SL6F9	B-1	0695h	1.50 GHz	600 MHz	400 MHz	Micro-FCPGA	2,5
SL6FA	B-1	0695h	1.60 GHz	600 MHz	400 MHz	Micro-FCPGA	2,5
SL6N5	B-1	0695h	1.70 GHz	600 MHz	400 MHz	Micro-FCPGA	2,5
SL6N8	B-1	0695h	1.30 GHz	600 MHz	400 MHz	Micro-FCBGA	2,6
SL6F5	B-1	0695h	1.40 GHz	600 MHz	400 MHz	Micro-FCBGA	2,5
SL6F6	B-1	0695h	1.50 GHz	600 MHz	400 MHz	Micro-FCBGA	2,5
SL6F7	B-1	0695h	1.60 GHz	600 MHz	400 MHz	Micro-FCBGA	2,5
SL6N9	B-1	0695h	1.70 GHz	600 MHz	400 MHz	Micro-FCBGA	2,5
SL6NC	B-1	0695h	1.10 GHz	600 MHz	400 MHz	Micro-FCBGA	1,2
SL6NB	B-1	0695h	1.20 GHz	600 MHz	400 MHz	Micro-FCBGA	1,2
SL6NA	B-1	0695h	1.30 GHz	600 MHz	400 MHz	Micro-FCBGA	1, 2
SL6NJ	B-1	0695h	900 MHz	600 MHz	400 MHz	Micro-FCBGA	3,4
SL6NH	B-1	0695h	1.00 GHz	600 MHz	400 MHz	Micro-FCBGA	3,4
SL6P4	B-1	0695h	1.10 GHz	600 MHz	400 MHz	Micro-FCBGA	3,4

NOTES:

1. VID[5:0] = 100001; VCC_CORE = 1.180 V for Highest Frequency Mode (HFM).
2. VID[5:0] = 101111; VCC_CORE = 0.956 V for Lowest Frequency Mode (LFM).
3. VID[5:0] = 101100; VCC_CORE = 1.004 V for Highest Frequency Mode (HFM).
4. VID[5:0] = 110110; VCC_CORE = 0.844 V for Lowest Frequency Mode (LFM).
5. VID[5:0] = 001110; VCC_CORE = 1.484 V for Highest Frequency Mode (HFM).
6. VID[5:0] = 010100; VCC_CORE = 1.388 V for Highest Frequency Mode (HFM).

Component Marking Information

Figure 1. Intel® Pentium® M Processor (Micro-FCPGA/FCBGA) Markings



§

Errata

Y1. Performance Monitoring Event That Counts Intel® Thermal Monitor 2 Transitions (59h) Is Not Accurate

Problem: The performance monitoring event that counts Intel Thermal Monitor 2 (Enhanced Intel SpeedStep® based) transitions may have inaccurate results.

Implication: There is no functional impact of this erratum. However this Performance Monitoring Event should not be used when accurate performance monitoring is required.

Workaround: None.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y2. Performance Monitoring Event that Counts the Number of Instructions Decoded (D0h) Is Not Accurate

Problem: The performance-monitoring event that counts the number of instructions decoded may have inaccurate results.

Implication: There is no functional impact of this erratum. However the results/counts from this Performance Monitoring Event should not be considered as being accurate

Workaround: None.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y3. RDTSC Instruction May Report the Wrong Time-stamp Counter Value

Problem: The Time-stamp Counter is a 64-bit counter that is read in two 32-bit chunks. The counter incorrectly advances and therefore the two chunks may go out of synchronization causing the Read Time-stamp Counter (RDTSC) instruction to report the wrong time-stamp counter value

Implication: This erratum may cause software to see the wrong representation of processor time and may result in unpredictable software operation.

Workaround: It is possible for BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y4. Code Segment Limit Violation May Occur on 4 Gbyte Limit Check

Problem: Code Segment limit violation may occur on 4 Gbyte limit check when the code stream wraps around in a way that one instruction ends at the last byte of the segment and the next instruction begins at 0x0.

Implication: This is a rare condition that may result in a system hang. Intel has not observed this erratum with any commercially available software, or system.

Workaround: Avoid code that wraps around segment limit.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y5. FST Instruction with Numeric and Null Segment Exceptions May Cause General Protection Faults to Be Missed and FP Linear Address (FLA) Mismatch

Problem: FST instruction combined with numeric and null segment exceptions may cause General Protection Faults to be missed and FP Linear Address (FLA) mismatch.

Implication: This is a rare condition that may result in a system hang. Intel has not observed this erratum with any commercially available software, or system.

Workaround: None.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y6. Code Segment (CS) Is Wrong on SMM Handler When SMBASE Is Not Aligned

Problem: With SMBASE being relocated to a non-aligned address, during SMM entry the CS can be improperly updated which can lead to an incorrect SMM handler.

Implication: This is a rare condition that may result in a system hang. Intel has not observed this erratum with any commercially available software, or system.

Workaround: Align SMBASE to 32K byte.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y7. IFU/BSU Deadlock May Cause System Hang

Problem: A lockable instruction with memory operand that spans across two pages may, given some rare internal conditions, hang the system.

Implication: When this erratum occurs, the system may hang. Intel has not observed this erratum with any commercially available software or system.

Workaround: Lockable data should always be contained in a single page.

Status: For the steppings affected, see the *Summary of Tables of Changes*.



Y8. Processor Can Enter a Livelock Condition under Certain Conditions When FP Exception Is Pending

Problem: Processor clock modulation may be controlled via a processor register (IA32_THERM_CONTROL) or via the STPCLK# signal. While the Processor clock is constantly being actively modulated at 12.5% and 25% duty cycles and there is a pending unmasked FP exception (ES pending), if you attempt a FP load (or Intel® MMX technology Mov instruction) and the load has an longer than typical latency the processor can enter a livelock.

Implication: When this erratum occurs, the processor will enter a livelock condition. Intel has not observed this erratum with any commercially available software or system.

Workaround: None.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y9. Write Cycle of Write Combining Memory Type Does Not Self Snoop

Problem: Write cycles of WC memory type do not self-snoop. This may result in data inconsistency- if the addresses of the WC data are aliased to WB memory type memory, which has been cached. In such a case, the internal caches will not be updated with the WC data sent on the system bus.

Implication: This condition may result in a data inconsistency. Intel has not observed this erratum with any commercially available software, system, nor components.

Workaround: Software should detect via the self-snoop bit in the CPUID features flags if the processor supports a self-snooping capability. Software should perform explicit memory management/flushing for aliased memory ranges on processors that do not self-snoop.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y10. Performance Monitoring Event that Counts Floating Point Computational Exceptions (11h) Is Not Accurate

Problem: Performance monitoring event that counts Floating Point Compare exceptions may have inaccurate results.

Implication: There is no functional impact of this erratum. However this Performance Monitoring Event should not be used when accurate performance monitoring is required.

Workaround: None.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y11. Inconsistent Reporting of Data Breakpoints on FP (Intel® MMX Technology) Loads

Problem: The reporting of data breakpoints on either FP or MMX technology loads is dependent upon the code faulting behavior prior to the execution of the load. If there is a fault pending prior to the execution of the load and FP exceptions are enabled there is a chance that data breakpoint on successive FP/MMX technology Loads may be reported twice.

Implication: Software debuggers should be aware of this possibility. There should be no implications to software operated outside of a debug environment.

Workaround: None.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y12. Code Breakpoint May Be Taken after POP SS Instruction If It Is followed by an Instruction That Faults

Problem: A POP SS instruction should inhibit all interrupts including Code Breakpoints until after execution of the following instruction. This allows sequential execution of POP SS and MOV ESP, EBP instructions without having an invalid stack during interrupt handling. However, a code breakpoint may be taken after POP SS if it is followed by an instruction that faults, this results in a code breakpoint being reported on an unexpected instruction boundary since both instructions should be atomic.

Implication: This can result in a mismatched Stack Segment and SP. Intel has not observed this erratum with any commercially available software, or system.

Workaround: As recommended in the IA32 Intel® Architecture Software Developer's Manual, the use "POP SS" in conjunction with "MOV ESP, EBP" will avoid the failure since the "Mov" will not fault.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y13. SysEnter and SysExit Instructions May Write Incorrect Requestor Privilege Level (RPL) in the FP Code Segment Selector (FCS)

Problem: SysEnter and SysExit instructions may write incorrect RPL in the FP Code Segment selector (FCS). As a result of this, the RPL field in FCS may be corrupted.

Implication: This is a rare condition that may result in a system hang. Intel has not observed this erratum with any commercially available software, or system.

Workaround: It is possible for BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary of Tables of Changes*.



Y14. Memory Aliasing with Inconsistent A and D Bits May Cause Processor Deadlock

Problem: In the event that software implements memory aliasing by having two Page Directory Entries (PDEs) point to a common Page Table Entry (PTE) and the Accessed and Dirty bits for the two PDEs are allowed to become inconsistent the processor may become deadlocked.

Implication: This erratum has not been observed with commercially available software.

Workaround: Software that needs to implement memory aliasing in this way should manage the consistency of the Accessed and Dirty bits.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y15. RDMSR or WRMSR to Invalid MSR Address May Not Cause GP Fault

Problem: The RDMSR and WRMSR instructions allow reading or writing of MSR's (Model Specific Registers) based on the index number placed in ECX. The processor should reject access to any reserved or unimplemented MSRs by generating #GP(0). However, there are some invalid MSR addressers for which the processor will not generate #GP(0). This erratum has not been observed with commercially available software.

Implication: For RDMSR, undefined values will be read into EDX:EAX. For WRMSR, undefined processor behavior may result.

Workaround: Do not use invalid MSR addresses with RDMSR or WRMSR.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y16. FP Tag Word Corruption

Problem: In some rare cases, fault information generated as the result of instruction execution may be incorrect. The result is an incorrect FP stack entry.

Implication: This erratum may result in corruption of the FP Tag Word in a way that a non-valid entry in the FP Stack may become valid. The software is not expected to read a non-valid entry. If the software attempts to use the stack entry (which is expected to be empty) the result may be an erroneous "Stack overflow".

Workaround: Do not disable SSE/SSE2 in control register CR4 and avoid Code segment Limit violation.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y17. Unable to Disable Reads/Writes to Performance Monitoring Related MSRs

Problem: The Performance Monitoring Available bit in the Miscellaneous Processor Features MSR (IA32_MISC_ENABLES.7) was defined so that when it is cleared to a 0, RDMSR/WRMSR/RDPMC instructions would return all zeros for reads of and prevent any writes to Performance Monitoring related MSRs. Currently it is possible to read from or write to Performance Monitoring related MSRs when the Performance Monitoring Available bit is cleared to a 0.

Implication: It is not possible to disallow reads and writes to the Performance Monitoring MSRs. Intel has not observed this erratum with any commercially available software or system.

Workaround: None.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y18. Move to Control Register Instruction May Generate a Breakpoint Report

Problem: A move (MOV) to Control Register (CR) instruction where Control Register is CR0, CR3 or CR4 may generate a breakpoint report.

Implication: MOV to Control Register Instruction is not expected to generate a breakpoint report.

Workaround: Ignore breakpoint data from MOV to CR instruction.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y19. REP MOVS Operation in Fast String Mode Continues in That Mode When Crossing into a Page with a Different Memory Type

Problem: A fast “REP MOVS” operation continues to be handled in fast mode when the string operation crosses a page boundary into an Uncacheable (UC) memory type. Also if the fast string operation crosses a page boundary into a WC memory region, the processor does not self snoop the WC memory region. This may result in incorrect data for the WC portion of the operation if those cache lines were previously cached as WB (through aliasing) and modified.

Implication: String elements should be handled by the processor at the native operand size in UC memory. In the event that the WB to WC aliasing case occurs, the end result could vary— from benign software execution to operating system or application failure. Intel has not observed either aspects of this erratum in commercially available software.

Workaround: Software operating within Intel’s recommendation will not require WB and WC memory aliased to the same physical address.

Status: For the steppings affected, see the *Summary Tables of Changes*.



Y20. The FXSAVE, STOS, or MOVS Instruction May Cause a Store Ordering Violation When Data Crosses a Page with a UC Memory Type

Problem: If the data from an FXSAVE, STOS, or MOVS instruction crosses a page boundary from WB to UC memory type and this instruction is immediately followed by a second instruction that also issues a store to memory, the final data stores from both instructions may occur in the wrong order.

Implication: The impact of this store ordering behavior may vary from normal software execution to potential software failure. Intel has not observed this erratum in commercially available software.

Workaround: FXSAVE, STOS, or MOVS data must not cross page boundary from WB to UC memory type.

Status: For the steppings affected, see the *Summary Tables of Changes*.

Y21. Machine Check Exception May Occur Due to Improper Line Eviction in the IFU

Problem: The processor is designed to signal an unrecoverable Machine Check Exception (MCE) as a consistency checking mechanism. Under a complex set of circumstances involving multiple speculative branches and memory accesses, there exists a one cycle long window in which the processor may signal a MCE in the Instruction Fetch Unit (IFU) because instructions previously decoded have been evicted from the IFU. The one cycle long window is opened when an opportunistic fetch receives a partial hit on a previously executed but not as yet completed store resident in the store buffer. The resulting partial hit erroneously causes the eviction of a line from the IFU at a time when the processor is expecting the line to still be present. If the MCE for this particular IFU event is disabled, execution will continue normally.

Implication: While this erratum may occur on a system with any number of processors, the probability of occurrence increases with the number of processors. If this erratum does occur, a machine check exception will result. Note systems that implement an operating system that does not enable the Machine Check Architecture will be completely unaffected by this erratum (e.g., Windows* 95 and Windows 98).

Workaround: It is possible for BIOS code to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y22. POPF and POPFD Instructions That Set the Trap Flag Bit May Cause Unpredictable Processor Behavior

Problem: In some rare cases, POPF and POPFD instructions that set the Trap Flag (TF) bit in the EFLAGS register (causing the processor to enter Single-Step mode) may cause unpredictable processor behavior.

Implication: Single-Step operation is typically enabled during software debug activities, not during normal system operation.

Workaround: There is no workaround for Single-Step operation in commercially available software. For debug activities on custom software the POPF and POPFD instructions could be immediately followed by a NOP instruction to facilitate correct execution.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y23. Performance Event Counter Returns Incorrect Value on L2_LINES_IN Event

Problem: The performance event counter returns an incorrect value on L2_LINES_IN event (EMON event #24H) when the L2 cache is disabled.

Implication: Due to this erratum, L2_LINES_IN performance event counter should not be monitored while the L2 cache is disabled. This erratum has no functional impact.

Workaround: Ignore L2_LINES_IN event when the L2 cache is disabled.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y24. VM Bit Will Be Cleared on a Double Fault Handler

Problem: Following a task switch to a Double Fault Handler that was initiated while the processor was in virtual-8086 (VM86) mode, the VM bit will be incorrectly cleared in EFLAGS.

Implication: When the OS recovers from the double fault handler, the processor will no longer be in VM86 mode

Workaround: None

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y25. Code Fetch Matching Disabled Debug Register May Cause Debug Exception

Problem: The bits L0-3 and G0-3 enable breakpoints local to a task and global to all tasks, respectively. If one of these bits is set, a breakpoint is enabled, corresponding to the addresses in the debug registers DR0-DR3. If at least one of these breakpoints is enabled, any of these registers are *disabled* (i.e., L_n and G_n are 0), and RW_n for the disabled register is 00 (indicating a breakpoint on instruction execution), normally an instruction fetch will not cause an instruction-breakpoint fault based on a match with the address in the disabled register(s). However, if the address in a disabled register matches the address of a code fetch which also results in a page fault, an instruction-breakpoint fault will occur.

Implication: While debugging software, extraneous instruction-breakpoint faults may be encountered if breakpoint registers are not cleared when they are disabled. Debug software which does not implement a code breakpoint handler will fail, if this occurs. If a handler is present, the fault will be serviced. Mixing data and code may exacerbate this problem by allowing disabled data breakpoint registers to break on an instruction fetch.

Workaround: The debug handler should clear breakpoint registers before they become disabled.

Status: For the steppings affected, see the *Summary of Tables of Changes*.



Y26. Upper Four PAT Entries Not Usable With Mode B or Mode C Paging

Problem: The Page Attribute Table (PAT) contains eight entries, which must all be initialized and considered when setting up memory types for the Pentium III processor. However, in Mode B or Mode C paging, the upper four entries do not function correctly for 4-Kbyte pages. Specifically, bit 7 of page table entries that translate addresses to 4-kbyte pages should be used as the upper bit of a 3-bit index to determine the PAT entry that specifies the memory type for the page. When Mode B (CR4.PSE = 1) and/or Mode C (CR4.PAE) are enabled, the processor forces this bit to zero when determining the memory type regardless of the value in the page table entry. The upper four entries of the PAT function correctly for 2-Mbyte and 4-Mbyte large pages (specified by bit 12 of the page directory entry for those translations).

Implication: Only the lower four PAT entries are useful for 4-kB translations when Mode B or C paging is used. In Mode A paging (4-Kbyte pages only), all eight entries may be used. All eight entries may be used for large pages in Mode B or C paging.

Workaround: None identified.

Status: For the steppings affected, see the *Summary of Tables of Changes*

Y27. SSE/SSE2 Streaming Store Resulting in a Self-Modifying Code (SMC) Event May Cause Unexpected Behavior

Problem: An SSE or SSE2 streaming store that results in a Self-Modifying Code (SMC) event may cause unexpected behavior. The SMC event occurs on a full address match of code contained in L1 cache.

Implication: Due to this erratum, any of the following events may occur:

1. A data access break point may be incorrectly reported on the instruction pointer (IP) just before the store instruction.
2. A non-cacheable store can appear twice on the external bus (the first time it will write only 8 bytes, the second time it will write the entire 16 bytes).

Intel has not observed this erratum with any commercially available software. This erratum has been seen in a synthetic test environment.

Workaround: None identified.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y28. Removed; See Erratum Y4

Y29. Removed; See Erratum Y5

Y30. Removed; See Erratum Y6

Y31. Page with PAT (Page Attribute Table) Set to USWC (Uncacheable Speculative Write Combine) While Associated MTRR (Memory Type Range Register) Is UC (Uncacheable) May Consolidate to UC

Problem: A page whose PAT memory type is USWC while the relevant MTRR memory type is UC, the consolidated memory type may be treated as UC (rather than WC as specified in IA-32 Intel® Architecture Software Developer's Manual).

Implication: When this erratum occurs, the memory page may be as UC (rather than WC). This may have a negative performance impact.

Workaround: None identified.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y32. Under Certain Conditions LTR (Load Task Register) Instruction May Result in System Hang

Problem: An LTR instruction may result in a system hang if all the following conditions are met:

1. Invalid data selector of the TR (Task Register) resulting with either #GP (General Protection Fault) or #NP (Segment Not Present Fault).
2. GDT (Global Descriptor Table) is not 8-bytes aligned.
3. Data BP (breakpoint) is set on cache line containing the descriptor data. When this erratum occurs, the memory page may be as UC (rather than WC). This may have a negative performance impact.

Implication: This erratum may result in system hang if all conditions have been met. This erratum has not been observed in commercial operating systems or software. For performance reasons, GDT is typically aligned to 8-bytes

Workaround: Software should align GDT to 8-bytes.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y33. Loading from Memory Type USWC (Uncacheable Speculative Write Combine) May Get Its Data Internally Forwarded from a Previous Pending Store

Problem: A load from memory type USWC may get its data internally forwarded from a pending store. As a result, the expected load may never be issued to the external bus.

Implication: When this erratum occurs, a USWC load request may be satisfied without being observed on the external bus. There are no known usage models where this behavior results in any negative side-effects

Workaround: Do not use memory type USWC for memory that has read side-effects.

Status: For the steppings affected, see the *Summary of Tables of Changes*.



Y34. FXSAVE after FNINIT without an Intervening FP (Floating Point) Instruction May Save Uninitialized Values for FDP (x87 FPU Instruction Operand (Data) Pointer Offset) and FDS (x87 FPU Instruction Operand (Data) Pointer Selector)

Problem: An FXSAVE after FNINIT without an intervening FP instruction may save uninitialized values for FDP and FDS.

Implication: When this erratum occurs, the values for FDP/FDS in the FXSAVE structure may appear to be random values. These values will be initialized by the first FP instruction executed after the FXRSTOR that restore the saved floating point state. Any FP instruction with memory operand will initialize FDP/FDS. Intel has not observed this erratum with any commercially available software.

Workaround: After an FNINIT, do not expect the FXSAVE memory image to be correct, until at least one FP instruction with a memory operand has been executed.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y35. FSTP (Floating Point Store) Instruction under Certain Conditions May Result In Erroneously Setting a Valid Bit on an FP (Floating Point) Stack Register

Problem: An FSTP instruction with a PDE/PTE (Page Directory Entry/Page Table Entry) A/D bit update followed by user mode access fault due to a code fetch to a page that has supervisor only access permission may result in erroneously setting a valid bit of an FP stack register. The FP top of stack pointer is unchanged.

Implication: This erratum may cause an unexpected stack overflow.

Workaround: User mode code should not count on being able to recover from illegal accesses to memory regions protected with supervisor only access when using FP instructions.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y36. Snoops during the Execution of a HLT (Halt) Instruction May Lead to Unpredictable System Behavior

Problem: If during the execution of a HLT instruction an external snoop causes an eviction from the instruction fetch unit (IFU) instruction cache, the processor may, on exit from the HLT state, erroneously read stale data from the victim cache.

Implication: This erratum may lead to unpredictable system behavior. Intel has only observed this condition in non-mobile configurations.

Workaround: It is possible for the BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y37. Invalid Entries in Page-Directory-Pointer-Table Register (PDPTR) May Cause General Protection (#GP) Exception If the Reserved Bits Are Set to One

Problem: Invalid entries in the Page-Directory-Pointer-Table Register (PDPTR) that have the reserved bits set to one may cause a General Protection (#GP) exception.

Implication: Intel has not observed this erratum with any commercially available software.

Workaround: Do not set the reserved bits to one when PDPTR entries are invalid.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y38. INIT Does Not Clear Global Entries in the TLB

Problem: INIT may not flush a TLB entry when:

1. The processor is in protected mode with paging enabled and the page global enable flag is set (PGE bit of CR4 register)
2. G bit for the page table entry is set
3. TLB entry is present in TLB when INIT occurs

Implication: Software may encounter unexpected page fault or incorrect address translation due to a TLB entry erroneously left in TLB after INIT.

Workaround: Write to CR3, CR4 or CR0 registers before writing to memory early in BIOS code to clear all the global entries from TLB.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y39. Use of Memory Aliasing with Inconsistent Memory Type May Cause System Hang

Problem: Software that implements memory aliasing by having more than one linear addresses mapped to the same physical page with different cache types may cause the system to hang. This would occur if one of the addresses is non-cacheable used in code segment and the other a cacheable address. If the cacheable address finds its way in instruction cache, and non-cacheable address is fetched in IFU, the processor may invalidate the non-cacheable address from the fetch unit. Any micro-architectural event that causes instruction restart will expect this instruction to still be in fetch unit and lack of it will cause system hang.

Implication: This erratum has not been observed with commercially available software.

Workaround: Although it is possible to have a single physical page mapped by two different linear addresses with different memory types, Intel has strongly discouraged this practice as it may lead to undefined results. Software that needs to implement memory aliasing should manage the memory type consistency.

Status: For the steppings affected, see the *Summary of Tables of Changes*.



Y40. Machine Check Exception May Occur When Interleaving Code between Different Memory Types

Problem: A small window of opportunity exists where code fetches interleaved between different memory types may cause a machine check exception. A complex set of micro-architectural boundary conditions is required to expose this window.

Implication: Interleaved instruction fetches between different memory types may result in a machine check exception. The system may hang if machine check exceptions are disabled. Intel has not observed the occurrence of this erratum while running commercially available applications or operating systems.

Workaround: Software can avoid this erratum by placing a serializing instruction between code fetches between different memory types.

Status: For the steppings affected, see the *Summary of Tables of Changes*.

Y41. Split I/O Writes Adjacent to Retry of APIC End of Interrupt (EOI) Request May Cause Livelock Condition

Problem: When Split I/O instruction writes occur adjacent to a retry of a Local APIC End of Interrupt (EOI) request by the chipset, a livelock condition may result. The required sequences of events are:

1. The processor issues a Local APIC EOI message.
2. The chipset responds with a retry because its downstream ports are full. It expects the processor to return with the same EOI request.
3. The processor issues a Split I/O write instruction instead.
4. The chipset responds with a retry because it expected the APIC EOI.
5. The processor insists the Split I/O write instruction must be completed and issues write instruction again.

Implication: A processor livelock may occur causing a system hang. This issue has only been observed in synthetic lab testing conditions and has not been seen in any commercially available applications. The erratum does not occur with Intel mobile chipset based platforms.

Workaround: Use the PIC instead of the APIC for the interrupt controller.

§

Specification Changes

There are no specification changes in this Specification Update revision.

§

Specification Clarifications

Y1. Specification Clarification with Respect to Time-stamp Counter

In the “Debugging and Performance Monitoring” section (Sections 15.8, 15.10.9 and 15.10.9.3) of the *IA-32 Intel® Architecture Software Developer’s Manual Volume 3: System Programming Guide*, the Time-stamp Counter definition has been updated to include support for the future processors. This change will be incorporated in the next revision of the *IA-32 Intel® Architecture Software Developer’s Manual*.

15.8 Time-Stamp Counter

The IA-32 architecture (beginning with the Pentium processor) defines a time-stamp counter mechanism that can be used to monitor and identify the relative time occurrence of processor events. The counter’s architecture includes the following components:

- **TSC flag** — A feature bit that indicates the availability of the time-stamp counter. The counter is available in an IA-32 processor implementation if the function CPUID.1:EDX.TSC[bit 4] = 1.
- **IA32_TIME_STAMP_COUNTER MSR** (called TSC MSR in P6 family and Pentium processors) — The MSR used as the counter.
- **RDTSC instruction** — An instruction used to read the time-stamp counter.
- **TSD flag** — A control register flag is used to enable or disable the time-stamp counter (enabled if CR4.TSD[bit 2] = 1).

The time-stamp counter (as implemented in the P6 family, Pentium, Pentium M, Pentium 4, and Intel Xeon processors) is a 64-bit counter that is set to 0 following a RESET of the processor. Following a RESET, the counter will increment even when the processor is halted by the HLT instruction or the external STPCLK# pin. Note that the assertion of the external DPSLP# pin may cause the time-stamp counter to stop.

Members of the processor families increment the time-stamp counter differently:

- For Pentium M processors (family [06H], models [09H, 0DH]); for Pentium 4 processors, Intel Xeon processors (family [0FH], models [00H, 01H, or 02H]); and for P6 family processors: the time-stamp counter increments with every internal processor clock cycle. The internal processor clock cycle is determined by the current core-clock to bus-clock ratio. Intel® SpeedStep® technology transitions may also impact the processor clock.
- For Pentium 4 processors, Intel Xeon processors (family [0FH], models [03H and higher]): the time-stamp counter increments at a constant rate. That rate may be set by the maximum core-clock to bus-clock ratio of the processor or may be set by the frequency at which the processor is booted. The specific processor configuration determines the behavior. Constant TSC behavior ensures that the duration of each clock tick is uniform and supports the use of the TSC as a wall clock timer even if the processor core changes frequency. This is the architectural behavior moving forward.

Note: To determine average processor clock frequency, Intel recommends the use of Performance Monitoring logic to count processor core clocks over the period of time for which the average is required. See Section 15.10.9 and Appendix A in this manual for more information.

The RDTSC instruction reads the time-stamp counter and is guaranteed to return a monotonically increasing unique value whenever executed, except for a 64-bit counter wraparound. Intel guarantees that the time-stamp counter will not wraparound within 10 years after being reset. The period for counter wrap is longer for Pentium 4, Intel Xeon, P6 family, and Pentium processors.

Normally, the RDTSC instruction can be executed by programs and procedures running at any privilege level and in virtual-8086 mode. The TSD flag allows use of this instruction to be restricted to programs and procedures running at privilege level 0. A secure operating system would set the TSD flag during system initialization to disable user access to the time-stamp counter. An operating system that disables user access to the time-stamp counter should emulate the instruction through a user-accessible programming interface.

The RDTSC instruction is not serializing or ordered with other instructions. It does not necessarily wait until all previous instructions have been executed before reading the counter. Similarly, subsequent instructions may begin execution before the RDTSC instruction operation is performed.

The RDMSR and WRMSR instructions read and write the time-stamp counter, treating the time-stamp counter as an ordinary MSR (address 10H). In the Pentium 4, Intel Xeon, and P6 family processors, all 64-bits of the time-stamp counter are read using RDMSR (just as with RDTSC). When WRMSR is used to write the time-stamp counter on processors before family [0FH], models [03H, 04H]: only the low order 32-bits of the time-stamp counter can be written (the high-order 32 bits are cleared to 0). For family [0FH], models [03H, 04H]: all 64 bits are writeable.

15.10.9 Counting Clocks

The count of cycles, also known as clockticks, forms the basis for measuring how long a program takes to execute. Clockticks are also used as part of efficiency ratios like cycles per instruction (CPI). Processor clocks may stop ticking under circumstances like the following:

- The processor is halted when there is nothing for the CPU to do. For example, the processor may halt to save power while the computer is servicing an I/O request. When Hyper-Threading Technology is enabled, both logical processors must be halted for performance-monitoring counters to be powered down.
- The processor is asleep as a result of being halted or because of a power-management scheme. There are different levels of sleep. In the some deep sleep levels, the time-stamp counter stops counting.

There are three ways to count processor clock cycles to monitor performance. These are:

- **Non-halted clockticks** — Measures clock cycles in which the specified logical processor is not halted and is not in any power-saving state. When Hyper-Threading Technology is enabled, these ticks can be measured on a per-logical-processor basis.
- **Non-sleep clockticks** — Measures clock cycles in which the specified physical processor is not in a sleep mode or in a power-saving state. These ticks cannot be measured on a logical-processor basis.
- **Time-stamp counter** — Some processor models permit clock cycles to be measured when the physical processor is not in deep sleep (by using the time-stamp counter and the RDTSC instruction). Note that such ticks cannot be measured on a per-logical-processor basis. See Section 10.8 for detail on processor capabilities.

The first two methods use performance counters and can be set up to cause an interrupt upon overflow (for sampling). They may also be useful where it is easier for a tool to read a performance counter than to use a time-stamp counter (the timestamp counter is accessed using the RDTSC instruction).



For applications with a significant amount of I/O, there are two ratios of interest:

- **Non-halted CPI** — Non-halted clockticks/instructions retired measures the CPI for phases where the CPU was being used. This ratio can be measured on a logical-processor basis when Hyper-Threading Technology is enabled.
- **Nominal CPI** — Time-stamp counter ticks/instructions retired measures the CPI over the duration of a program, including those periods when the machine halts while waiting for I/O.

15.10.9.3 Incrementing the Time-Stamp Counter

The time-stamp counter increments when the clock signal on the system bus is active and when the sleep pin is not asserted. The counter value can be read with the RDTSC instruction.

The time-stamp counter and the non-sleep clockticks count may not agree in all cases and for all processors. See Section 10.8 for more information on counter operation.

§

Documentation Changes

There are no documentation changes in this Specification Update revision.